



AI for HR and workforce capability

Using AI responsibly for learning, workforce support, and people operations.

WORKS WITH THE ASSISTANTS TEAMS ALREADY USE



EDITION 11 / EMBER & LILAC

EXECUTIVE WORKING EDITION - 08 OCT 2026

AT A GLANCE

Using AI responsibly for learning, workforce support, and people operations.

HR sits in two positions at once: it uses AI to serve employees and manage people processes, and it is responsible for building the whole workforce's ability to use AI well. Both roles carry unusual responsibility, because decisions about hiring, pay, performance, and learning affect people's livelihoods and are often subject to employment, anti-discrimination, and data protection rules. This playbook covers employee service, role-based learning, workforce planning, and the limits that keep people-related AI fair, explainable, and under human judgment.

WHO IT IS FOR

Chief people officers, heads of HR and learning and development, HR operations and HR business partner leaders, talent and workforce planning teams, and employee relations and works council contacts.

TIME TO USE

60 minutes to read; one working session to apply

What you will have at the end

- A map of HR use cases split into low-risk employee support and high-stakes people decisions
- A role-based capability plan that links AI skills to actual jobs rather than generic courses
- Guardrails for fairness, transparency, and employee data, with named human decision points
- A 90-day plan combining an employee service pilot with a workforce capability baseline

Contents

01	Two jobs for HR: use AI well and build a workforce that can	03
02	Use cases, fit, and boundaries	04
03	How people-facing AI works, and why capability is not a course	06
04	Roles, governance, and employee voice	08
05	Fairness, privacy, transparency, and trust	09
06	Baseline, metrics, and capability evidence	11
07	A practical rollout for HR and capability	12
08	Common failures, an illustrative scenario, and vendor questions	14
09	A first quarter that produces evidence	15
10	Terms used in this publication	16

EXECUTIVE BRIEF

Two jobs for HR: use AI well and build a workforce that can

01



HR is both a user of AI in people processes and the owner of workforce capability. Doing the second well matters more to the business than doing the first quickly.

Most HR conversations about AI begin with recruiting tools and policy chatbots. These matter, but they are the smaller half of the task. The larger half is helping every function change how it works: which roles will use AI daily, which skills they need, how performance expectations shift, and how people are supported while that happens. Organizations that treat AI capability as a one-time training event find that a few enthusiasts use the tools well while most people avoid them or use them badly.

The people-process half carries distinct risk. When AI helps answer a benefits question, an error is annoying. When it helps screen candidates, rank staff, or flag attrition risk, errors can be discriminatory, hard to explain, and in some jurisdictions regulated. The sound approach is to separate employee support from people decisions, apply light governance to the first, and apply strict human judgment, testing, and transparency to the second.



Four areas of work

- Employee service: answering policy and process questions and helping managers with routine people tasks
- Role-based learning: building the skills each job family needs, with practice on real work
- Workforce planning: understanding which tasks change, which roles evolve, and where capability gaps will appear
- People operations: drafting job descriptions, summarizing feedback, preparing communications, and supporting case handling



DECISIONS ABOUT PEOPLE STAY WITH PEOPLE

AI may inform, prepare, and summarize. A named manager or HR professional must make and own decisions on hiring, promotion, pay, discipline, and termination, and must be able to explain the reasons without citing a model score.

WORKING QUESTION

Which of your current AI conversations in HR is about tools, and which is about the capability your people need?

WHERE IT FITS

Use cases, fit, and boundaries

02

Sort HR uses by the consequences for an individual. The greater the effect on someone's job or pay, the stronger the human role and the testing must be.

A helpful sorting device is to ask what happens to a person if the output is wrong. If an employee gets an unclear answer about parental leave and asks again, the cost is small. If a screening tool quietly filters out qualified candidates from a group, the cost is large, hard to detect, and possibly unlawful. This difference should drive how much testing, oversight, and transparency each use case receives.

USE CASE	FIT	MAIN RISK AND CONTROL
Policy and benefits questions for employees	Strong	Wrong or stale policy. Ground in current policy, cite it, route complex cases to HR

USE CASE	FIT	MAIN RISK AND CONTROL
Drafting job descriptions, communications, and learning content	Strong	Biased wording or errors. HR reviews; use inclusive language checks
Role-based learning plans and practice scenarios	Strong	Generic content. Build from real tasks and manager input
Summarizing engagement survey comments	Moderate	Loss of nuance and anonymity. Protect small groups; human reads themes
Workforce planning scenarios	Moderate	Overconfident forecasts. Treat as scenarios with stated assumptions
Candidate screening, ranking, interview analysis	High risk	Bias, explainability, legal exposure. Strict testing, human decisions, counsel review
Performance scoring, attrition prediction, monitoring	Avoid by default	Fairness, trust, and legality concerns. Only with clear purpose, consultation, and legal advice

Examples of boundaries to write down

- No decision about hiring, pay, promotion, discipline, or termination is made or finalized by a model
- No inference of protected characteristics, health, or emotional state from text, voice, or video
- No covert monitoring of employees with AI; any monitoring has a stated purpose and is disclosed
- Sensitive case notes, such as grievances and medical information, do not enter general-purpose tools
- Employees are told when they interact with an automated assistant and how to reach a person

EMPLOYMENT AND AI RULES ARE TIGHTENING

Employment law, anti-discrimination law, data protection law, works council or consultation rules, and the EU AI Act (which treats some employment uses as high risk) may apply, depending on where you operate. Check current requirements with employment counsel before using AI in hiring or evaluation.

WORKING QUESTION

Which HR task in your function, if wrong for one employee, would you be unable to explain to them face to face?

HOW IT WORKS

How people-facing AI works, and why capability is not a course

03

Employee assistants answer from governed policy content. Capability programs work when they build habits on real tasks, not when they deliver tool tours.

An HR assistant typically retrieves the relevant policy passages and asks a language model to answer using only those. It can be useful for routine questions, such as how to book leave or what the expenses limit is. It fails when policies are out of date, inconsistent across countries, or written in dense legal language. As with customer service, the first deliverable of an AI project is often a cleanup of policy documents, with owners, effective dates, and country variants clearly marked.

01 Identify the question

The employee asks in everyday language. The system checks whether the question is in scope and whether it touches a sensitive topic.

02 Retrieve current policy

It finds the right policy and country variant, and where allowed the employee's own relevant data, such as leave balance.

03 Answer with a source

The assistant replies in plain language and links to the policy. It states when it is unsure.

04 Route when needed

Complex, sensitive, or emotional queries go to an HR professional with the conversation attached.

05 Learn from gaps

Unanswered and poorly rated questions go to policy owners as a queue of fixes.

What makes capability building work

Generic AI courses produce awareness but little change. Capability grows when learning is tied to a role and a task: an account manager practices building a client brief, a finance analyst practices drafting variance commentary and checking it, a recruiter practices writing inclusive job ads and spotting biased output. Each person learns what the tool can do for their work, what it gets wrong, how to check it, and what data they must not enter. Managers need their own track, since they set expectations and judge quality.

- Start with job families and the specific tasks in each that AI changes
- Use real documents and scenarios, with practice followed by on-the-job application
- Teach checking habits: verify facts, test edge cases, never paste restricted data
- Create a network of local champions who share tested examples, not hype
- Refresh content frequently, since tools and policies change

WORKING QUESTION

If you asked ten employees what they may and may not do with AI at work, how many would give the same answer?

OPERATING MODEL

Roles, governance, and employee voice

04

People-related AI needs joint ownership between HR, legal, IT, and the business, with employees and their representatives informed from the start.

HR cannot own this alone. Technology owns systems and security. Legal and data protection own compliance. Business leaders own how work changes. HR owns policy, fairness, capability, and the relationship with employees. A small cross-functional group, with the authority to approve use cases and pause them, keeps decisions consistent and prevents each function from buying its own tool. Include employee voices through representatives, works councils where they exist, or pilot user panels.

ROLE	OWNS	DECIDES
Chief people officer	People strategy, risk appetite	Which HR uses are permitted and which are off limits
HR operations and policy owners	Policy content, employee service quality	Accuracy standards and escalation to humans
Learning and development	Capability framework and programs	Role profiles, curricula, and assessment of skills
Business leaders and managers	How work and roles change	Task redesign, expectations, and coaching
Legal and employment counsel	Compliance with employment and AI rules	Approval of high-risk uses and notices
Data protection officer and security	Employee data handling	Data categories, retention, and vendor approval
Employee representatives or panels	Employee perspective	Consulted on monitoring, scope, and communication

A capability framework in practice

Define three or four levels of AI capability that map to real expectations. For example: aware (understands what AI is, the rules, and the risks), practitioner (uses approved tools effectively in their role and checks output), advanced practitioner (designs workflows and prompts others use), and steward (owns governance for a team or process). Assess each job family against these levels, identify the gaps, and plan training and practice accordingly. Avoid turning this into a score for individuals that feeds into pay or promotion without careful thought, since doing so changes behavior and can feel punitive.

BE HONEST ABOUT JOB CHANGE

Employees will ask whether AI threatens their job. Do not promise what you cannot guarantee. Explain which tasks are changing, what support exists for reskilling and redeployment, and how decisions will be made fairly. Silence is read as bad news.

WORKING QUESTION

Who in your organization would be accountable if an HR tool produced a biased result, and do they know it?

RISKS AND CONTROLS

Fairness, privacy,
transparency, and trust

05

The risks specific to HR are discriminatory outcomes, misuse of employee data, opaque decisions, and loss of trust.

Models learn patterns from historical data and language, which can include historical bias. A tool trained or prompted on past hiring decisions may reproduce the preferences of those decisions. Even apparently neutral inputs can act as proxies for protected characteristics. This is why high-stakes people tools demand testing for differences in outcomes across groups, documented reasons for the factors used, and a human who can override.

RISK	EXAMPLE	CONTROL
Discrimination	Screening favors some groups without a job-related reason	Bias testing, job-related criteria, human decision, counsel review
Employee privacy	Sensitive case notes or health data entered into a general tool	Data classification, approved tools, minimization, strict access
Opacity	Employee cannot learn why a decision was made	Explainable criteria, records of reasoning, right to ask for review
Wrong policy advice	Assistant gives an outdated entitlement	Governed content, citations, owner review cycle
Surveillance creep	Productivity or sentiment monitoring introduced quietly	Purpose limits, disclosure, consultation, legal review
Anonymity loss	Survey comments traced to small teams	Aggregation thresholds, human review of themes
Over-reliance by managers	Drafted feedback or ratings accepted unread	Manager training, examples, and review of samples

Before using AI in any process that affects individuals

- The purpose is specific, documented, and job-related
- A lawful basis and any required notices or consultations have been checked with counsel
- The tool has been tested for outcome differences across relevant groups using suitable data
- A named human makes and owns the final decision and can override the tool
- People affected can ask for an explanation and a human review
- Data retention, access, and vendor terms have been reviewed by data protection

Rules about automated decision-making, candidate and employee notices, worker consultation, and bias audits differ widely and are changing. General frameworks such as the NIST AI Risk Management Framework and ISO/IEC 42001 can structure your governance, but they do not replace legal advice on the rules where you hire and employ.

WORKING QUESTION

For your most consequential people decision, can you show that a human genuinely weighed the evidence and was free to disagree with the tool?

MEASURING VALUE

Baseline, metrics, and capability evidence

06



Measure employee outcomes and capability, not only usage. High adoption of a poor habit is not a success.

For employee service, measure resolution and accuracy, not just question counts. For capability, measure changes in how work is done and in results, not course completions. A useful baseline for capability is a short task-based assessment for each job family, run before and after training: for example, ask a sample of people to complete a realistic task with AI and score the quality, speed, and checking behavior. Combine this with manager observations and adoption data from approved tools.

METRIC	HOW TO MEASURE	WATCH FOR
HR query resolution and accuracy	Sampled answers graded against current policy; repeat questions	Counting questions answered instead of questions answered correctly
Time to answer and HR case volume	Median response time; cases per employee by type	Cases shifting rather than disappearing
Employee satisfaction with HR service	Short post-interaction rating, segmented	Low response from dissatisfied staff
Task-based capability score	Same realistic task before and after training, scored by rubric	Training to the test
Sustained, appropriate tool use	Active use of approved tools by role plus sampled quality	Use of unapproved tools; usage without value
Time-to-productivity and quality in target roles	Ramp time and quality measures for new starters or role changes	Confounding changes in the business
Fairness indicators in high-stakes tools	Outcome rates by group, and override rates by humans	Small samples; conclusions without expert statistical help



Build a test set for employee service

Collect 75 to 150 real employee questions, anonymized, across countries and topics, and write the correct answer for each, including questions that should be escalated. Run the assistant

against the set before launch and after each policy or prompt change. Grade accuracy, completeness, tone, and correct routing. Include tricky cases, such as questions about leave during a disciplinary process or entitlements that vary by country. Choose your pass mark with HR and legal, for instance no incorrect statements of entitlement and all sensitive cases routed to a person.



ILLUSTRATIVE EXAMPLE

An HR assistant shows high usage in its first month. Sampling finds that for part-time staff it quotes full-time leave rules. Usage hid an accuracy gap that only the graded sample found. The team adds worker-type variants to policies and adds part-time cases to the test set.



WORKING QUESTION

What evidence would convince you that a training program changed how people work, rather than what they say about AI?



IMPLEMENTATION GUIDE

A practical rollout for HR and capability

07

Run two tracks in parallel: a low-risk employee service pilot and a role-based capability program, both backed by a clear policy on what is off limits.



01 Set the rules first

Publish an acceptable-use policy covering approved tools, data restrictions, disclosure, and the list of HR uses that are off limits or high risk.

02 Fix the policy content

Consolidate and date HR policies, mark country and worker-type variants, and assign an owner to each. This is the foundation for any employee assistant.

03 Pilot employee service

Launch an assistant for a narrow set of topics, such as leave and expenses, to a pilot group, with easy routing to HR and weekly review of gaps.

04 Map roles to capability

Identify the job families most affected, define the AI tasks in each, and set capability levels and baseline assessments.

05 Train on real work

Deliver short, role-based sessions using live tasks, followed by practice and manager coaching. Appoint champions in each function.

06 Review and adjust

Examine results quarterly, update policy and curricula, and decide which further uses, if any, meet the bar for fairness and legal review.

Readiness checklist

- An acceptable-use policy is published and understood by managers
- Employee representatives or a panel have been informed or consulted as required
- Policy documents have owners, dates, and variants marked
- Escalation to a person works and is visible to employees
- Data protection has approved the data flows and vendors
- The capability framework links to actual tasks in each job family
- Treat managers as a priority audience; their behavior sets what employees believe is acceptable
- Offer learning time during working hours, since optional evening courses reach only the already enthusiastic
- Share good examples across functions so that people see what is possible and safe
- Plan for reskilling and redeployment alongside any efficiency goals
- Communicate regularly and plainly, including about what has not been decided

WORKING QUESTION

Which job family would benefit most from a practical, task-based AI program in the next 90 days?

PITFALLS AND BUYING

Common failures, an illustrative scenario, and vendor questions

08

Failures typically come from starting with high-stakes tools, treating training as a one-off, or leaving policy content in disarray.

- Buying a screening or ranking tool before defining job-related criteria and how bias will be tested
- Deploying an assistant over contradictory policies and then blaming the model
- Measuring completions of a generic course and calling it capability
- Allowing unapproved tools to spread because the approved route is slow or unclear
- Announcing AI efficiency goals without a credible plan for people, which erodes trust

ILLUSTRATIVE SCENARIO, NOT A REAL COMPANY

A manufacturer trials a candidate ranking tool. Recruiters start trusting the order of the list. A review finds that the tool favors candidates whose resumes resemble past hires, and few recruiters opened lower-ranked profiles. The company stops using ranking, uses AI for summarizing and scheduling only, and keeps screening decisions with trained recruiters.

Questions to ask any vendor

Vendor diligence

- What evidence exists of testing for bias, and can we run our own tests on our data
- How are employee and candidate data stored, used, retained, and deleted, and where
- Does the vendor use our data to train shared models
- Can the tool explain which factors influenced an output in terms a person can understand
- What human override and audit logging does it provide
- Which jurisdictions' employment and AI rules has the vendor assessed, and what notices does it support
- What happens to our data and configurations if we end the contract

Ask for a pilot on your own data with a defined test, and have employment counsel and data protection review the results before any wider rollout. Be most skeptical of claims that a tool can assess personality, emotion, or potential from video, voice, or text; these claims are often contested and may be restricted by law in some places.



WORKING QUESTION

Which of your planned HR uses would you be uncomfortable explaining publicly, and what does that tell you?

30 / 60 / 90

A first quarter that produces evidence

09



Set the rules, fix the policy content, pilot employee service, and build capability by role, with strict human oversight where people are affected.



First 30 days: ground it

- Publish an acceptable-use policy with approved tools, data rules, and off-limits HR uses
- Inventory and date HR policies, mark variants, and assign owners
- Inform employee representatives or a panel about the plan
- Define three or four AI capability levels and select priority job families

Days 31 to 60: test it

- Pilot an employee assistant on a narrow topic set with easy routing to HR
- Build and run a test set of 75 to 150 anonymized employee questions
- Run baseline task-based capability assessments for the priority job families
- Start role-based sessions with managers and local champions



Days 61 to 90: decide

- Review accuracy, satisfaction, and policy gaps from the pilot and decide on scope
- Reassess capability against the baseline and adjust curricula
- Set the governance rhythm and criteria for any high-stakes HR use, including counsel review
- Publish a plain-language update to employees on what has changed and what comes next

GLOSSARY

Terms used in this publication

10

TERM	MEANING
Acceptable-use policy	A document setting out which AI tools may be used, with what data, and for which purposes.
Adverse impact	A pattern in which a process disadvantages a protected group, even without intent, which testing aims to detect.
Automated decision-making	Decisions made with little or no human involvement, which are regulated in some places when they significantly affect individuals.
Capability framework	A structured description of the skills and proficiency levels expected in roles or job families.
Champion network	Respected practitioners in each function who share tested practice and help colleagues adopt new ways of working.
Data minimization	Collecting and using only the personal data needed for a defined purpose.
Employee data	Personal information about staff, some of which, such as health or disciplinary records, is especially sensitive.
Explainability	The ability to describe in understandable terms why a system or person reached a result.
Job family	A group of roles with related skills and tasks, used for planning and learning.
Proxy variable	An apparently neutral factor that correlates with a protected characteristic and can cause indirect bias.

TERM	MEANING
Reskilling	Training people in new skills so they can move into changed or different roles.
Workforce planning	Analyzing future demand for work and skills against current capacity to guide hiring, development, and redeployment.

A note on use

Free assessments are available at halden.company/assessments, and you can reach the team at halden.company/contact. Figures here are illustrative, and nothing in this playbook is legal advice. Employment and AI rules vary by location; consult qualified counsel.

TAKE THE NEXT STEP

Take the free Halden assessment at halden.company/assessments, or talk to us at halden.company/contact about turning this into a plan for your organization.